Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning
Main Track

General Game Playing With State-Independent Communication

Sean Zammit and **Michael Thielscher**
UNSW Sydney, Australia
z5207033@zmail.unsw.edu.au, mit@unsw.edu.au

Abstract

Communication actions in games are usually given meaning by either the effect they have on the game state or a possible reduction in the size of the information set of an agent. However, this precludes analysis of games involving *state-independent communication*, in which players are given the ability to communicate with each other but are not required to be truthful. As such, these communication actions cannot be used to reduce the size of the information set or update beliefs in states without considering the intent of the communicating agent. In this paper, we introduce a language to describe the rules of such games as an extension of the *Game Description Language (GDL)*. We also identify a set of scenarios involving state-independent communication actions in which an effective agent should be able to derive information, and propose and evaluate strategies for reasoning about such actions in these scenarios.

1 Introduction

General game players are programs designed to play arbitrary games, the rules of which they are provided with at runtime. Initially proposed by Barney Pell in a paper first published in 1992 and republished in 1996 (Pell 1996), general game players must be designed to play a broad range of unknown games without relying on the individual features of particular games. This requirement to function effectively in a range of possible scenarios makes the field of General Game Playing relevant to the broader context of general AI research (Genesereth and Thielscher 2014).

Our focus is on a specific type of game that has gone unanalysed in the context of General Game Playing — games of state-independent communication. In these games, players are given the ability to communicate with each other, but unlike in the majority of research on communication games, players are not required to be truthful. This presents novel challenges for general game-playing systems in deciding what claims to make to other players, and in gaining information from others' claims despite their potential lack of truth. In this way, these games are more akin to how real-world communication functions, with no physical law requiring certain conditions to be true in order to make a given statement. Current approaches to communication games have several limitations, especially in regard to state-independent forms of communication, which can be summarised as follows.

Limitation 1. *In the context of General Game Playing, there is no useful language for formally defining games of state-independent communication.*

Limitation 2. *Existing analyses of games which typically involve state-independent communication simplify their considerations by making all communication state-dependent, or omitting communication entirely.*

Limitation 3. *There is limited strategic analysis of the effects of state-independent communication or the effect of permitting lies in probabilistic belief models.*

These limitations motivate the work described in this paper. We will begin by addressing Limitation 1 by introducing a language capable of representing meaningful state-independent communication. There are several languages available for describing games of incomplete information, such as the *Game Description Language* GDL-II (Thielscher 2010), or as of more recently Ludii (Piette et al. 2019). While we need to choose a specific language to extend, we aim for the concepts used to be as general as possible such that it would be simple to apply the extension we describe to other languages in the future if desired. As such, we prefer to base our extension on a language like GDL with a simple set of keywords in order to not make our extension heavily reliant on the particular features of an individual language. We will demonstrate that our extended language GDL-II$^{\mathrm{C}}$ is capable of expressing a wide variety of state-independent communication actions. We will then use this language to describe and identify strategies for state-independent communication scenarios and analyse the effectiveness of our proposed strategies, addressing Limitations 2 and 3.

The remainder of the paper is organised as follows. We first provide a background on general game playing with incomplete-information games in Section 2, including a definition for state-independent communication, and a brief review of the existing literature on such games. Section 3 presents syntax and semantics of our new general game description language for describing, and playing, such games. Section 4 identifies a set of common communication scenarios and proposes belief revision strategies an agent could use in each scenario. Section 5 analyses the effectiveness of these strategies with and against agents using other strategies. Finally, Section 6 evaluates the contributions presented and identifies potential avenues for future work.

2 Background

2.1 General Game Playing

GDL-II In General Game Playing, agents are given descriptions of the game they are to play at runtime using a known language for describing game rules, rather than being pre-programmed or trained already knowing the precise details of a game. It expresses game rules in the form of a logic program similar to Prolog. A small set of predefined keywords are sufficiently expressive to describe a range of incomplete information games:

`role``(R)`	`R` is a player
`init``(F)`	`F` holds in the initial position
`true``(F)`	`F` holds in the current position
`legal``(R,M)`	`R` can do move `M` in the current position
`does``(R,M)`	player `R` does move `M`
`next``(F)`	`F` holds in the next position
`terminal`	the current position is terminal
`goal``(R,N)`	`R` gets `N` points in the current position
`sees``(R,P)`	`R` perceives `P` in the next position
`random`	the random player

The first eight keywords come from standard GDL (Love et al. 2006), while the final two were added to support the description of games of incomplete information; after each action, players receive all of their *percepts* as defined by the rules for the **`sees`** keyword, while the **`random`** keyword provides the ability to describe events such as cards drawn from a deck of cards (Schiffel and Thielscher 2015).

Communication games We define *communication games* to be games of incomplete or imperfect information in which the communication of information between players is a strategic consideration. This definition encompasses *hidden-role games* such as Werewolf (see below), in which the 'role' and objectives of each player are determined randomly from a set of possible allocations, with the majority of players not knowing the roles of most of the other players, and *cooperative communication games* such as Hanabi, in which players share an objective, but each have knowledge about different parts of the game state. There are also many games that incorporate elements of one or both of these categories in addition to other non-communication elements.

State-independent communication We define actions as being *state-independent communication* actions when two criteria are met:

1. The choice of which action to take is based on what information a player wishes to communicate.
2. The action is legal in all states that are possible from the perspective of any player, given any knowledge they have of the game rules and past actions.

Importantly, state-independent communication actions do not provide the receivers with knowledge about the game state that must be true simply because the action to communicate it was taken. Rather, just like regular communication in a real-world environment, receivers draw meaning with which to inform beliefs, making decisions about how strongly they trust the claim being made or making inferences about other properties of the game state.

Werewolf A particularly-researched hidden-role game is Werewolf, in which players are assigned the roles of either 'wolf' or 'villager' at the beginning of the game. Wolves know who the other wolves are, but only specific villagers know the roles of other players and so their faction has a numeric advantage to compensate. At night, the wolves kill a villager, while during the day all players discuss and vote on a player to kill, with day-night cycles repeating until only one faction remains.

The day phase is of particular interest to us, when players are free to make arbitrary claims about the identities or special roles of themselves or other players. This arbitrariness is missing from the analysis of other papers. Girlea, Amit, and Girju (2014), for example, define a series of axioms for specific communication actions from which meaning can be derived. One such axiom states that if a player claims to be the 'seer' (a role which knows the identity of other players), then they must either be the seer or a wolf. This makes the claim of being the seer a state-*dependent* communication action — it limits the possible states to only those in which the player is actually the seer or a wolf, and not another role.

2.2 Limitations of Existing Languages

As noted, the representation of Werewolf by Girlea et al. expressed communication in a state-dependent manner. For specific games, re-expressing communication as being state-independent is a fairly simple matter of defining some communication protocol. For Werewolf this has been done by the AI Wolf Contest (Osawa 2022), both in an explicitly defined protocol-based division, and also in a natural language division — while natural language might not be specifically designed for Werewolf, it is a system of terms with agreed-upon meanings through which state-independent communication may be performed. These protocols have enabled the work of Dolça (2019) and Shoji et al. (2019), which correlate a player's number of utterances with their win-rate.

In General Game Playing, however, such a protocol does not exist. In fact, in existing languages such as GDL-II, it is not possible to define state-independent communication actions that carry inherent meaning. Consider the Cooperative Spy game by Schofield and Thielscher (2015): Two spies are disarming a bomb. The first knows which wire to cut and may communicate that to the other spy, while the second has the ability to cut a wire. Both win if the correct wire is cut, otherwise both lose. There appears to be a simple strategy for each agent — to be truthful and trusting. However, the difficulty is apparent if we try to represent communication in the game rules. Consider the below definition for communication of which wire to cut in GDL-II:

```
legal(spy1, tell_which_wire(W)) :-
    true(correct_wire(W))
sees(spy2, wire_to_cut(W)) :-
    does(spy1, tell_which_wire(W))
```

This definition is an example of state-dependent communication. Meaning is given by a reduction in the set of possible game states when the first spy performs the action of telling the second spy which wire to cut, as the only way to legally perform that action is for the specified wire to be the correct wire to cut. However, when we modify this definition to be state-independent, we encounter an issue:

```
legal(spy1, tell_which_wire(W)) :-
    is_a_wire(W)
sees(spy2, wire_to_cut(W)) :-
    does(spy1, tell_which_wire(W))
```

Now, even though the first spy can legally say any wire is the correct wire to cut, its statements have no inherent meaning. There is nothing in the game rules to suggest that the second spy seeing `wire_to_cut(red)` as a result of the first spy's communication should infer that `correct_wire(red)`. This problem occurs for all attempts to describe state-independent communication in GDL-II. Because meaning is given by reducing the set of possible game states to only those where the action could have been taken, any communication actions defined in this way are by definition state-dependent, and any state-independent communication actions have no inherent meaning. To the best of our knowledge, this limitation applies to all existing languages used to formally define general games, as none of them have been designed with the goal of enabling state-independent communication in mind.

2.3 On Communication Strategies

Dynamic epistemic logics (DEL) for modelling knowledge change after truthful communications are well-studied and provide insight into what players know, what they know others know and so on, before and after players receive communications or take actions that affect their knowledge of the world (van Ditmarsch, Der, and Koi 2008), including in the context of general game playing based on GDL (Engesser et al. 2021). However, methods of updating agents' models of the world after potentially false, state-independent communication are less well-examined.

One such analysis is that of van Ditmarsch (2013), who builds on the *consecutive numbers riddle* (van Emde, Groenendijk, and Stokhof 1984): Anne and Bill are secretly told a positive number. Their numbers will be one apart, and they are aware of all of these details. The following truthful conversation now takes place:

- Anne: “I do not know your number”
- Bill: “I do not know your number”
- Anne: “I know your number”
- Bill: “I know your number”

Suppose Anne is told 2 and Bill was told 3. Similarly to games such as *muddy children* (van Ditmarsch, Der, and Koi 2008), announcements that an agent does not know their number are actually informative to other agents: Anne, who has a 2, knows that Bill has a 1 or a 3. Bill, who has a 3, knows that Anne has a 2 or a 4. For Anne to know Bill's number, she would need to have the number 1, forcing Bill to have 2. By announcing she does not know Bill's number, she tells Bill she does not have a 1 (which in this case is not new information for Bill). However, then Bill announces he does not know Anne's number. By the reasoning above, this rules out Bill having a 1, which reduces the possible states from Anne's perspective from 2 to 1.

Van Ditmarsch then analyses several scenarios in which Anne and Bill are not perfectly honest. In the first scenario, Anne lies in her first announcement. Bill cannot reconcile her announcement with his existing knowledge that she couldn't know his number yet, so he can immediately conclude Anne was lying. In the second scenario, Bill lies in his first announcement. If Anne believes him, she will conclude he must have had a 1, and so she announces she knows his number. While incorrect, if Bill believes *she* is telling the truth, he can now conclude that she must have had a 2, and so announce he knows her number.

There are two important details to observe here though: First, Anne may not be perfectly trusting and truthful. Second, in the current version of the game, neither have a motivation to lie. Van Ditmarsch (2013) describes how the first observation leads into a couple of different options for Anne. If she instead distrusts Bill's lie, she might claim that he is lying, but be uncertain of that fact. She could also announce that she doesn't know his number. However, at this point Bill observes that Anne is either lying or doesn't believe he is telling the truth.

The second observation is particularly important for our own design and analysis of games and strategies later in this paper. There's no benefit to either agent to lie, in fact it is in each player's best interests to cooperate with the other. If the players can figure out from the form of the game that cooperation is in their best interests, they should be able to achieve a high *belief* that they know their numbers.

3 A Language to Describe Games of State-Independent Communication

In this section we present our extended general game description language *GDL-II with Claims* (GDL-II^{C}) for describing games that include state-independent communication, that is, actions which have meaning outside of a restriction of the set of possible game states. Our goal is to be able to describe a variety of types of games of state-independent communication and at the same time use the minimal number of keywords sufficient for achieving this.

3.1 Formal Language Specification

We extend GDL-II by the ability to describe ‘claims’. A claim is a communication *percept*, rather than an action, which we can use to imbue state-independent communication actions with meaning.

Syntax The syntax of GDL-II^{C} is the same as GDL-II, with the addition of a new keyword, **`claim`**`(S,P)`. The **`claim`** keyword may only appear as the percept in the head of a **`sees`** clause. Its two parameters denote the identity `S` of the player “sending” the claim, and information `P` claimed by that player to be true in the state the claim was sent in. Parameter `P` should be in the form of an atom used in the GDL game description, such as **`true`**`(is_wolf(p3))` or `and(a,b)`, where `and` is derived from other facts. Syntactically, `P` may contain terms that do not exist elsewhere in the game description, such as if `is_wolf` was not described. Such claims could never hold in any state and become essentially meaningless.

Semantics As we observed in Section 2.2, any attempt to describe state-independent communication in existing languages results in communication actions without inherent meaning. The solution is to define a protocol known to all agents that can be used to convey meaning in communications. In our case, we want to enable individual game descriptions to each define their own set of permissible claims to form a communication protocol for their described game. We achieve this by extending the definition of the semantics of a game in GDL-II by a *claim relation* as follows.

Definition 1. *Let G by a valid GDL-II*$^{\mathrm{C}}$ *specification whose signature determines the set of ground terms* Σ. *The semantics of G is the game* $(R, s_1, t, l, u, \mathcal{I}, g, \pi, \mathcal{C})$ *where the first eight terms are given by the same assignments as in GDL-II:*[1]

- $R \subseteq \Sigma$ *(the roles);*
- $s_1 \subseteq \Sigma$ *(the initial position);*
- $t \subseteq 2^\Sigma$ *(the terminal positions);*
- $l \subseteq R \times \Sigma \times 2^\Sigma$ *(the legality relation);*
- $u : (R \mapsto \Sigma) \times 2^\Sigma \mapsto 2^\Sigma$ *(the update function);*
- $\mathcal{I} \subseteq R \times (R \mapsto \Sigma) \times 2^\Sigma \times \Sigma$ *(the information relation);*
- $g \subseteq R \times \mathbb{N} \times 2^\Sigma$ *(the goal relation);*
- $\pi : (R \setminus \{\mathtt{random}\} \mapsto \Sigma) \times 2^\Sigma \mapsto \mathcal{P}(2^\Sigma)$*;*

and the new claim relation

$$\mathcal{C} \subseteq R \times R \times (R \mapsto \Sigma) \times 2^\Sigma \times \Sigma$$

is determined from the game rules G as follows:

$$\begin{aligned}\mathcal{C} \;=\; \{\; & (s, r, M, S, p) : G \cup M^{\mathtt{does}} \cup S^{\mathtt{true}} \\ & \vdash \mathtt{sees}(r, \mathtt{claim}(s, p)), \\ & s, r \in R,\; M : (R \mapsto \Sigma),\; S \in 2^\Sigma,\; p \in \Sigma\}.\end{aligned}$$

The meaning of the above assignment is that the claim relation $\mathcal{C}$ contains all claims that can possibly be made in the game, represented by 5-tuples denoting for every

- sender s and receiver r,
- set of clauses $S^{\mathtt{true}}$ denoting a current state, and
- joint move of all players $M^{\mathtt{does}}$

each fact p being claimed to r by s when move $M^{\mathtt{does}}$ is taken in state $S^{\mathtt{true}}$. Based on the semantics of a game, we can extend the formal design of a *Game Manager* to games with state-independent communication described in GDL-II$^{\mathrm{C}}$ as follows:

1. Send each $r \in R \setminus \{\mathtt{random}\}$ the GDL-II$^{\mathrm{C}}$ rules and inform them about their individual roles r. Set $S := s_1$.
2. Collect the individual moves from each player. Depending on the context, this may be after a set amount of time allotted for agents to consider their moves. For the random player, choose with uniform probability an element from the set $\{m : (\mathtt{random}, m, S) \in l\}$ as their move. Set M to be the joint move of all players.
3. Send to each $r \in R \setminus \{\mathtt{random}\}$ the set of percepts $\{p : (r, M, S, p) \in \mathcal{I}\}$ and the set of claimed information $\{(s, p) : (s, r, M, S, p) \in \mathcal{C}\}$. Set $S := u(M, S)$.
4. Repeat steps 2 and 3 until the state is terminal, i.e. $S \in t$. Determine the result n for $r \in R \setminus \{\mathtt{random}\}$ by $(r, n, S) \in g$.

Steps 1, 2 and 4 are taken from the original implementation of a Game Manager for GDL-II (Thielscher 2010). Step 3 is similar, but adds to the data sent to each player the set of claimed information they receive when a move occurs.

Interpreting of the meaning of claims It is tempting to try to concretely define the meaning of a claim in terms of the intents or beliefs of agents, but ultimately, we have to recognise that claims only have meaning through the interpretations of agents. There is nothing preventing two agents from deciding to interpret any protocol's equivalent of "Cutting the red wire defuses the bomb" as meaning that the blue wire is the correct wire to cut — instead, it relies on accepted conventions for conveying meaning, just as natural language does. What we *do* provide by defining claims in the way that we have is a convention for meaning that may be used by General Game Playing agents across all games, and a method for enabling communication to occur.

3.2 Expressiveness

To demonstrate that GDL-II$^{\mathrm{C}}$ is capable of expressing a wide variety of state-independent communication actions, we provide examples in this subsection for some common use-cases, using the game of Werewolf as an illustrative example.

Public communication In public communication, every agent receives the claim being made. We might describe public accusations that a player is a wolf with the following rules:

```
legal(S, accuse_as_wolf(W)) :-
    is_day_phase, role(S), role(W)
sees(R, claim(S, true(is_wolf(W)))) :-
    does(S, accuse_as_wolf(W)), role(R)
```

The first clause states that is is legal for any player S to take the action of accusing another player W of being a wolf as long as the game is currently in the day phase. The second clause uses the new **`claim`** keyword to specify that when that action is taken, each player R receives a claim from the sender S that **`true`**`(is_wolf(W))` held in the game state at the time that the claim was made.

Note that in this definition the sender also receives their own claim. However, because each receiver is given the identity of the sender along with the claimed information, they can discard any claims made by themselves if they so choose.

Private communication In private communication, not all agents receive the claim being made. The following is an example of how we might describe the rules for some player 'whispering' to another player a claim that some player is a

[1]For a full explanation of these existing components of the semantics of GDL-II, we refer the reader to the original definition (Thielscher 2010).

wolf in the game of Werewolf: [2]

```
legal(S, whisper_wolf_accusation(R, W)) :-
    is_day, role(S), role(R), role(W)
sees(R, claim(S, true(is_wolf(W)))) :-
    does(S, whisper_wolf_accusation(R, W))
```

Similarly to our example for public accusations, our first clause defines that it is legal to make a private accusation that a player W is a wolf to another player R in the day phase. Our second clause then uses the receiver specified by the action from the first clause to restrict the receivers of the claim to only the player R.

We can use similar techniques to our example for public communication if we want claims to be visible to some group of players. By the following rule we can permit wolves to make claims amongst themselves about which player they think is the seer, with their communication only seen by other wolves:

```
legal(S, wolf_whisper_about_seer(P)) :-
    is_night, true(is_wolf(S)), role(P)
sees(R, claim(S, true(is_seer(P)))) :-
    does(S, wolf_whisper_about_seer(P)),
    true(is_wolf(R))
```

Claims about historical state In some games, we want players to be able to make claims about something that was true in some previous game state. The following rules show how we might permit players to make such claims:

```
1  init(stored_history(is_alive(X))) :-
2      role(X)
3  ...
4  next(stored_history(P)) :-
5      true(stored_history(P))
6  next(stored_history(prev(P))) :-
7      true(stored_history(P))
8  next(history(prev(P))) :-
9      true(P), true(stored_history(P))
10 next(history(prev(P))) :-
11     true(history(P))
12 legal(S, claim_history(prev(P))) :-
13     stored_history(prev(P)), role(S)
14 sees(R, claim(S, true(history(P)))) :-
15     does(S, claim_history(P)), role(R)
```

We begin by specifying each of the facts we want to permit players to make claims about; in this case we use as an arbitrary example the ability to make Lines 3–6 update the `stored_history` predicate each turn such that the initially specified facts, and their previous states only back until the initial state are known to be stored in the history. Rule 7–8 adds each currently true fact that can be stored in the history to the history, while rule 9–10 shifts each element in the history each turn to refer to one state further back in time. This completes the construction of our history, and lines 11-14 enable agents to make claims about the history.

The key insight here is keeping track of what *could* be in the history to restrict the set of possible claims by agents, preserving *allowed*-ness of the game description as specified by Lloyd and Topor (1986). Having done so, we can treat claims about the history just like any other claims.

Claims with logical connectives Finally, we demonstrate how we can permit players to make claims in the form 'A and B' and 'A or B':

```
can_connect(is_alive(X)) :- role(X)
...
and(P, Q) :- true(P), true(Q),
    can_connect(P), can_connect(Q)
legal(S, claim_both(P, Q)) :- role(S),
    can_connect(P), can_connect(Q)
sees(R, claim(S, and(P, Q))) :-
    does(S, claim_both(P, Q)), role(R)
or(P, Q) :- true(P),
    can_connect(P), can_connect(Q)
or(P, Q) :- true(Q),
    can_connect(P), can_connect(Q)
legal(S, claim_either(P, Q)) :- role(S),
    can_connect(P), can_connect(Q)
sees(R, claim(S, or(P, Q))) :-
    does(S, claim_either(P, Q)), role(R)
```

4 Strategies for Communication Scenarios

In this section we propose and justify strategies for a number of typical scenarios that occur in games of state-independent communication. We also present simple games in which each scenario occurs for evaluating the ability of agents to play state-independent communication games.[3]

4.1 General Scenarios

First, we examine a category of scenarios that occur generally, rather than due to the particular composition of a group of players and their objectives.

A Claim Advantageous Regardless of Truth

Scenario: A player makes a claim that they would benefit from the listener believing regardless of whether or not the claim is true.

Strategy: When such a claim is made, the listener should not change their beliefs. If the player making the claim has no incentive to tell the truth, then no information can be confidently gained from their claim. In GGP, we can consider having 'an incentive to tell the truth' to be equivalent to the agent expecting a higher utility from telling the truth than lying. As the way agents calculate expected utility may vary by strategy, one agent may believe that there is an incentive to tell the truth while another may not.

Game — Good or Evil: Consider a game with two players, a guesser and a subject. The subject is randomly assigned the role of either 'good' or 'evil'. The guesser's goal is to correctly guess which role the subject has been assigned, while the subject has the goal of the guesser deciding that they are 'good', regardless of whether this is true. Additionally, the subject is able to claim to be either of the two roles.

[2] Note that typically private communication is not permitted in the Werewolf game; we choose to use it here as an example for easy comparison with our example for public communication

[3] The descriptions for these games in GDL-II^{C} are provided at https://github.com/z5207033/ZT23-Game-Descriptions

In this game, both a 'good' and 'evil' subject would want the guesser, or listener, to believe that they are 'good', and have the ability to make such a claim. There is no reason for the listener to believe them though — both types of subjects would act identically if they thought they would be believed.

A Claim Disadvantageous if False

Scenario: A player makes a claim that would be detrimental to them if the listener believed it and it were not true.

Strategy: Increase belief in the possible game states in which the claim holds to a near-certainty. The player would not have rationally made such a claim if it were false.

Game — Cooperative Spies: Recall the Cooperative Spy game (Schofield and Thielscher 2015) described in Section 2.2. Two spies are disarming a bomb where they have to correctly choose which of two wires to cut. The first spy knows which wire to cut, and may tell the other spy that either wire is the correct one, while that spy chooses which wire to cut.

In this game, the first spy has an incentive to not lie. If they lied and were believed, both spies would lose.

4.2 Group Scenarios

The scenarios of this second category occur depending on the particular combinations of goals for a set of players.

A Claim Between Allied Players

Scenario: Two players want the same outcomes, and one makes a claim to the other.

Strategy: The player making the claim may have an incentive to lie to their partner if they are limited in the amount of information they can share, and their partner believing them will lead to them taking the most advantageous actions. Because of this, the receiver of the claim should believe the other player, who would rationally be acting in the receiver's best interests.

Game — Shared Envelopes: Consider a two-player game with two envelopes, each with some random amount of money from \$0 to \$100. Similarly to the Cooperative Spy game, the first player knows how much money is in each envelope, and the second chooses an envelope. Both players get the payoff of however much money is in the envelope. However, the only action the first player is permitted to make is a single claim about how much money is in just one of the envelopes.

The first player should claim that \$100 is in whichever envelope has the most money, in order to persuade the second agent to pick that envelope even if the amount of money in it is small. The second agent should increase belief in the possible game states in which the claim holds to a near-certainty, even knowing that the first agent may have an incentive to lie.

A Public Claim

Scenario: In a game with at least 3 players and public claims, player 1 has no incentive to deceive player 2, and does have an incentive to deceive player 3. Player 2 knows this. Player 1 makes a public claim heard by all players.

Strategy: The listeners should consider whether to believe the claim based on what beliefs would be beneficial to the claiming player in each of the possible game states. Player 2, who knows that player 1 benefits from deceiving player 3, should reduce their belief in states where player 1's claim is true. Player 3, whose information set may encompass states where player 1 has an incentive to lie and states where they should tell the truth, should reduce belief in states where player 1 has an incentive to lie and their claim is true and states where they have an incentive to tell the truth and their claim is false.

Game — Probable Shared Envelopes: Consider a similar game to Shared Envelopes, this time with 3 players and 2 envelopes, one containing \$0 and the other containing \$100. The game has a 90% chance of being of a form where all 3 players get the payoff from the selected envelope, with the contents known by player 1 and the envelope chosen by player 3. It has a 10% chance of being in a form where the player 3 gets the payoff from their selected envelope, player 1 gets the payoff from the other envelope, and player 2 gets the payoff from an envelope they select (which may be the same envelope as player 3). Only players 1 and 2 know which form the game is in, and player 1 may claim either envelope has \$100.

This design creates a scenario in which player 3 should usually believe player 1's claim, but in which player 1 in the second form of the game has an incentive to deceive player 3, and player 2 knows they have this incentive.

For player 3, there are 4 possible initial states — form 1 with \$100 in the first or second envelope, each with 45% likelihood, and form 2 with \$100 in the first or second envelope, each with 5% likelihood. Suppose without loss of generality that player 1 claims the money is in the first envelope. Player 3 should reduce their beliefs in form 1 with the money in the second envelope and form 2 with the money in the first envelope to almost 0, as player 1's claim would not be rational with these initial state. This leaves form 1 with the money in the first envelope with a probability of $\approx 90\%$, and form 2 with the money in the second envelope with a probability of $\approx 10\%$, and so player 3 should choose the first envelope.

Performing similar considerations, in the first form player 2 makes no decisions, so we consider the second form, which they know the game is in. By the same logic as above, when player 1 claims the money is in the first envelope, player 2 reduces their belief in the game state where that claim holds to almost 0, and so selects the second envelope.

A Claim Contradicting Another Agent

Scenario: In a game with at least 3 agents, player 1 observes players 2 and 3 making contradictory claims.

Strategy: Player 1 should increase their belief in states where each player had an incentive to claim what they did, rather than believing either specific claim.

Game — One-Shot Investigation: Consider a game with three players, an investigator and two suspects. Each suspect independently has a 50% chance of being guilty — it is possible both or neither are guilty. The suspects also know whether they or the other suspect is guilty, and can make claims about the guilt of both suspects. The investigator has to guess which, if any, suspects are guilty. The investigator and innocent suspects win if the investigator guesses correctly, while guilty suspects win if *no* guilty suspect is guessed. Note that if there are two guilty suspects and the investigator only guesses one, or one guilty suspect and the investigator guesses both, all players lose.

While if both suspects are allied neither has an incentive to accuse the other, our scenario occurs when one suspect is guilty and the other innocent. The innocent suspect has an incentive to accuse the guilty suspect, as they need the investigator to guess that suspect is guilty in order to win. The innocent suspect has no reason to say this if the accused suspect was innocent, so the accused suspect has an incentive to counter-accuse, resulting in conflicting claims.

The investigator's strategy should be to reduce belief in the game states where the two players are allied to almost 0, as they would have reason to not contradict each other in these states. This leaves their beliefs in the game states where each suspect is the sole guilty player both at $\approx 50\%$.

4.3 A General Strategy for Receiving Claims

The specific strategies outlined above stem from the same general idea: Rather than choosing whether to believe a claim made by a player, receivers should modify their beliefs based on which states it would have made sense for the player to have made that claim. One way to do this, which we will evaluate in Section 5, is for a receiver to determine for each state in their own information set what the information set and beliefs of the sender would be if they followed this general strategy. The receiver can then reduce belief in states where the player making the claim would have from their perspective a lower expected utility if their claim were believed than if they were to make some other action (communication or otherwise). If this is the case in all possible states, the receiver defaults to believing the claim.

4.4 A General Strategy for Making Claims

Given the above strategy for receiving claims, we can define a general strategy for when to make claims:

1. Assume that the other players will follow the belief revision strategy described above and act to maximise their expected utility according to the probabilities assigned to each belief.
2. Treat the decision to make state-independent communication actions in the same way as one would consider regular actions, and choose the action that maximises expected utility based on the behavioural expectations of the other agents.

The key feature of this strategy is that it brings state-independent communication actions in line with how agents consider the value of regular actions. This neatly combines the ability to evaluate state-independent communication actions with an agent's ability to examine regular actions, so existing techniques such as Information Set variants of MCTS (Cowling, Powley, and Whitehouse 2012) can be applied in a state-independent communication context without modification to evaluate the utility of actions in an information set tree that is too large to feasibly fully search. Additionally, it enables agents to compare the expected utilities of state-independent communication and regular actions in game states which have both types as legal moves.

5 Evaluation of Strategies

5.1 Implementation

In order to demonstrate the efficacy of our strategies in the context of General Game Playing, we empirically tested agents using our strategies against agents using several other general strategies for each of the games we presented in Section 4. Our testing system was modelled off our design for a GDL-II^{C} Game Manager and provided agents with the same information and prompts that would be expected in a real system.[4]

Player strategies

There are 5 player implementations we will compare:

1. **Random**: A random player who chooses with uniform probability one of their legal moves to play.
2. **Regular**: Representing a typical general game playing agent, this player maximises their expected utility against agents using the same strategy as itself, but assigns no meaning to state-independent communication actions. In most cases this player functions as a Random agent and is grouped under **Random*** where applicable in our results due to the lack of information available to it.
3. **Truthful**: A player who randomly selects a truthful legal claim to make.
4. **Trusting**: A player based on the Regular player, but who trusts every claim made. If trusting every claim results in an empty set of possible states, this player falls back on making a random legal move.
5. **Proposed**: A player using the general strategies described in Section 4. Like the Regular player, it maximises its expected utility against agents using the same strategy as itself, but incorporates probabilistic beliefs by assigning a weighting to each game state in its information set that can be modified as determined by our strategies. In our implementation, 'believing' a state to be possible or not possible is achieved by multiplying or dividing respectively the weight for that state by 1000, and so this player can never have a belief of 0 in all states.

In our simple games, a player will only ever make or receive claims, and not both. The Truthful player is thus omitted as a strategy for players receiving claims, and the Trusting player is omitted as a maker of claims.

[4]The source code and setup for our experiments is provided at https://github.com/z5207033/ZT23-Experimental-Setup

Claim Maker		Random*	Truthful	Proposed
Guesser	Random*	50, 50	50, 50	50, 50
Good or	Trusting	50, 50	100, 50	50, 100
Evil Game	Proposed	50, 25	75, 25	50, 50
Guesser	Random	50, 50	50, 50	50, 50
Weighted	Regular	75, 0	75, 0	75, 0
Good or	Trusting	50, 50	100, 25	50, 50
Evil Game	Proposed	75, 0	75, 0	75, 0

Table 1: Oppositional average game results[5]

Claim Maker		Random*	Truthful	Proposed
Guesser	Random*	50, 50	50, 50	50, 50
Cooperative	Trusting	50, 50	100, 100	100, 100
Spies	Proposed	50, 50	100, 100	100, 100
Guesser	Random*	50, 50	50, 50	50, 50
Shared	Trusting	50, 50	65, 65	70, 70
Envelopes	Proposed	50, 50	65, 65	70, 70

Table 2: Cooperative average game results

Player 1		Random*	Truthful	Proposed
Player 3	Random*	50, 50	50, 50	50, 50
90% allied,	Trusting	50, 50	100, 90	90, 100
10% opp.	Proposed	50, 50	100, 90	90, 100
	Random*	50, -	50, -	50, -
Player 2 *Knows form*	Trusting	50, -	100, -	0, -
	Proposed	50, -	0, -	100, -

Table 3: Probable Shared Envelopes average game results[6]

5.2 Results

1,000,000 matches were played of each game for each combination of player strategies, and informally verified mathematically to ensure correctness. As any difference in average utility between strategies was fairly large, results were rounded to the nearest whole number to make our results easier to read. Our goal was not to quantify the relative strengths of strategies, as such numbers were highly dependent on the design of each of our games, but rather to identify patterns in which strategies were the most effective at different games.

5.3 Analysis

For the games we analyse, player combinations involving the Random player provide us with a useful baseline to compare against. We can observe that due to the communicative nature of our games, if either the sender or receiver of a claim is the Random player, then in most cases the other player may as well be random — when the sender is random, the receiver doesn't have reliable information, and when the receiver is random the sender's strategy has no effect.

Oppositional games In the Good or Evil game described in Section 4, the Proposed player has the best results as the player making claims. It assumes that its opponent will trust it if it states that it is evil as if it received such a claim it would trust it and always guess that the sender was evil, and so it always claims to be good. This is ignored by a Random* guesser, and disregarded by the Proposed guesser, but has benefits against a Trusting guesser.

As the guesser, the Proposed strategy is either equal with or behind a Trusting strategy for each type of sending strategy. However, this is dependent on two factors we would not expect in most useful contexts:

1. The Trusting strategy is only better against a Truthful agent, which in our results is at most as effective as our Proposed strategy and usually strictly worse, so competent game players should not employ the Truthful strategy.
2. If we modify the game slightly to have the player making the claim have a 75% chance of being evil, the Trusting strategy becomes worse than both the Regular and Proposed strategies against non-Truthful strategies, as seen in the lower half of Table 1, as it more often trusts a deceptive agent.

We also observe that all strategies as the sender perform worse against the Proposed agent. This is due to it disregarding claims that are always in the sender's interests (Truthful and Proposed senders), while penalising claims that aren't in the sender's interests (Random and Truthful senders).

Cooperative games In the Cooperative Spies game, any combination of non-Random* agents achieves the maximum utility, shown in Table 2. This is because the strategies either cooperate by default (Truthful and Trusting) or can determine that cooperation is in both player's best interests in the case of the Proposed strategy.

In the more complex Shared Envelopes game[7] however, we can see that a simple Truthful strategy isn't necessarily the best strategy in all cooperative games. When there are low payoffs in both envelopes, telling the truth about how much money is in the envelope with the most money (which the Truthful strategy does 50% of the time) results in a Trusting or Proposed agent picking the other envelope, expecting it to likely contain more money. The Proposed sender instead claims that there is some amount of money more than $50 in the envelope with the most money, ensuring that a Trusting or Proposed agent selects that envelope.

Mixed games The remaining games we analysed contained a mix of oppositional and cooperative elements. The first of these was the Probable Shared Envelopes game, for which the results are shown in Table 3. Player 1, as the player making claims, is punished for being Truthful by the

[5]In these tables, each result lists the average rewards received by the players on the left and top of the table respectively.

[6]For simplicity, we only display the expected utilities of Player 2 in the oppositional form for the game, as they take no actions in the coooperative form.

[7]To reduce the size of the information set, for our testing we reduce the possible amounts in each envelope to one of $0, $25, $50, $75 or $100.

Suspects		Rand*, Rand*	Rand*, True	Rand*, Prop	True, True	True, Prop	Prop, Prop
Investigator	Random*	25, 31, 31	25, 31, 31	25, 31, 31	25, 31, 31	25, 31, 31	25, 31, 31
	Trusting	25, 31, 31	44, 36, 36	31, 36, 48	100, 50, 50	63, 38, 69	38, 69, 69
	Proposed	25, 19, 19	54, 24, 29	39, 20, 45	88, 38, 38	75, 25, 50	50, 50, 50

Table 4: One-Shot Investigation average game results

Trusting and Proposed strategies in the 10% of oppositional instances of the game, while the Proposed strategy can effectively determine when lying is beneficial.

The Trusting and Proposed strategies are equally effective for Player 3, as the Proposed strategy determines that most of the time it is in Player 1's best interests to tell the truth, and so trusts any claims made. Player 2's results are more interesting, as the Proposed strategy assumes Player 1 will act in their own best interests and so believes that Player 1 is lying to Player 3. As such, the Proposed strategy's effectiveness is dependent on a competent Player 1.

Finally, we have the One-Shot Investigation game, the results of which are shown in Table 4. There is a lot of data in this game's results, so we focus on several key trends.

First, we observe that in general, a Trusting investigator is more effective than a Random investigator, and a Proposed investigator is more effective again except in the case where all suspects are Truthful, as has been the case in previous games. Notably, however, the Proposed investigator is more effective in all cases where there is only one Truthful suspect — while the Trusting investigator can be confused by contradictory information from a Random or Proposed suspect, the Proposed investigator can gain enough information to rule out several states and gain an advantage.

Second, the Proposed strategy as a suspect is always at least as good as the Random or Truthful strategies, and is strictly better whenever the investigator is not Random. This, we expect, is due to the Proposed strategy's ability to handle both cooperative and oppositional scenarios, both of which appear in this game depending on the allocations of guilt to suspects. By contrast, the Truthful strategy has only matched the effectiveness of the Proposed strategy in previous games which were entirely cooperative, and is strictly worse than the Proposed strategy in oppositional and mixed scenarios.

6 Discussion

GDL-II$^{\mathbf{C}}$ We have defined a game description language for describing games with state-independent communication. Where previously strategies and players of such games were limited to the specific game they were designed for, this language makes it possible to formally define games of state-independent communication in a way that is able to be read and understood by General Game Playing agents, enabling the examination of these games in a GGP context.

Additionally, we have presented definitions for a range of types of state-independent communication in our language. This facilitates the creation of definitions for future games.

Finally, as our extension was designed to be simple and as unreliant on the specific features of the language we extended as possible, we expect that it can be easily adapted to be applied to other game description languages if desired.

Scenarios We identified a number of typical scenarios in games of state-independent communication and proposed simple games with which General Game Playing strategies can be evaluated. We believe the One-Shot Investigation game to be of particular significance due to the range of scenarios that can occur and the complexity of the 3-way interaction between agents, while having a small and easily searchable information set tree.

Strategies Our results demonstrate that in nearly all cases, our Proposed strategy is at least as, and usually more effective than other simple strategies.

In oppositional games, the Proposed strategy outperformed the Random strategy as a receiver against all strategies except occasionally itself. It was worse than the Trusting strategy only when the sender was Truthful, which would not be a strategy chosen by capable players.

In purely cooperative games, the Proposed strategy was equivalent to a Trusting strategy as a receiver, as it was capable of determining that the other player had an incentive to tell the truth. As a sender, it matched or outperformed all other strategies by being able to lie when such an action was in the best interests of the group.

In mixed games, our Proposed strategy as a receiver was particularly effective due to its ability to more accurately determine when the sender had an incentive to lie, outperforming all other strategies except in oppositional cases where all senders were irrationally Truthful. As a sender, it was also successful, outperforming all other strategies in all games except when the receiving player determining the game outcome was the Random player.

As such, if we make the assumption then that well-designed agents will not use the strategy of always being truthful, we can expect a Proposed receiver to be a consistently more effective strategy than the other simple strategies we have examined.

Future work We have mostly focused on cooperative and hidden-role games in this paper, but could extend our analysis to games using state-independent communication for bluffs and negotiation in addition to actions on a shared game board, such as Sheriff of Nottingham.

Our strategies could also be combined with heuristic search techniques such as Information Set variants of MCTS in order to be applied to games with information set trees that are too large to fully search, or augmented with other strategies that update beliefs based on the communication actions that players choose *not* to make.

References

Cowling, P.I.; Powley, E.J.; and Whitehouse, D. 2012. Information set Monte Carlo tree search. *IEEE Transactions on Computational Intelligence and AI in Games* 4(2):120–143.

van Ditmarsch, H; Der, V.; and Kooi, B. 2008. *Dynamic Epistemic Logic*. Dordrecht: Springer.

van Ditmarsch, H. 2013. Dynamics of lying. *Synthese* 191(5):745–777.

Dolça, T. 2019. Are talkative AI agents more likely to win the Werewolf game? In *Proceedings of the 1st International Workshop of AI Werewolf and Dialog System*, 11–14.

van Emde Boas, P; Groenendijk, J.; and Stokhof, M. 1984. The Conway paradox: Its solution in an epistemic framework. In *Truth, Interpretation and Information: Selected Papers from the Third Amsterdam Colloquium*, 159–182.

Engesser, T.; Mattmüller, R.; Nebel, B; and Thielscher, M. 2021. Game Description Language and Dynamic Epistemic Logic compared. *Artificial Intelligence* 292:103433.

Girlea, C.; Amir, E.; and Girju, R. 2014. Tracking beliefs and intentions in the Werewolf game. In *Proceedings of the International Conference on Knowledge Representation and Reasoning*, 638–641.

Genesereth, M.R. and Thielscher, M. 2014. *General Game Playing*. San Rafael, Ca: Morgan & Claypool.

Love, N.; Hinrichs, T.; Haley, D.; Schkufza, E.; and Genesereth, M. 2006. *General Game Playing: Game Description Language Specification*. Tech. rep. LG-2006-01, Stanford Logic Group, Computer Science Department, Stanford University, 353 Serra Mall, Stanford, CA 94305. Available at: games.stanford.edu.

Lloyd, J.W. and Topor, R.W. 1986. A basis for deductive database systems II. *The Journal of Logic Programming* 3(1):55–67.

Osawa, H. *AIWolf contest*. [online] Available at: http://aiwolf.org/en/aiwolf_contest [Accessed 1 Dec. 2022].

Pell, B. 1996. A strategic metagame player for general chess-like games. *Computational Intelligence* 12(1):177–198.

Piette, E.; Soemers, D.J.; Stephenson, M.; Sironi, C.F.; Winands, M.H.; and Browne, C. 2019. Ludii–the ludemic general game system. *arXiv preprint arXiv:1905.05013*.

Schofield, M. and Thielscher, M. 2015. Lifting model sampling for general game playing to incomplete-information models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3585–3591.

Schiffel, S. and Thielscher, M. 2014. Representing and reasoning about the rules of general games with incomplete information. *Journal of Artificial Intelligence Research*, 49:171–206.

Shoji, N.; Jotaro, A.; Kosuke, O.; Kotaro, S.; Hideyuki, S.; Tatsunori, M.; and Noriko, K. 2019. Strategies for an autonomous agent playing the 'Werewolf game' as a stealth werewolf. In *Proceedings of the 1st International Workshop of AI Werewolf and Dialog System*.

Thielscher, M. 2010. A general game description language for incomplete information games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 994–999.